Dr. Adel Saleh
Program Manager
Advanced Technology Office

Next-Generation Global DoD Enterprise Network

In addition to all their other great accomplishments, the ancient Egyptians, my ancestors, were a famously seafaring people, but, all they had to build their ships were twisted acacia trees, barely taller than a man. So, they learned to build ships with short slats of wood, about 3 feet in length. Well, centuries pass, and their travels bring them to Lebanon, where the majestic cedar trees grew to 100 feet or more.

The Egyptians were stunned. They ordered their crews to fell the cedars, after which, they cut them into 3-foot pieces and built the ships the way they always had. Eventually, the Egyptians came to see the potential of long-plank design. But, this took some time.

If we, today, want to realize the compelling vision of network-centric warfare, we cannot afford to wait to incorporate the most advanced technology and the latest thinking available to us into developing the thing that makes it all possible: the network itself. That is to say, the fibers and switches that enable these powerful applications we're talking about.

Next-Generation Global DoD Enterprise Network

Let us get more specific, the Global DoD Enterprise Network forms the backbone of the DoD Global Information Grid (GIG). It consists of satellite-based networks, an optical core network, optical edge networks, and land-based fixed wireless networks, all with strong connectivity to tactical wireless networks. Incidentally, all this is virtually separate from the public internet, but, today shares some of its resources.

My focus is on the ground-based, next generation optical core and edge networks. What do I mean by "core" and "edge"? As an analogy, the core network is like the interstate highway system, and the edge networks are like local roadways.

The core network, which has to have a global reach, uses high-speed links, e.g., fibers carrying a large number of wavelengths. It directly interconnects, say, 100 or fewer core nodes, where traffic is added, dropped, or switched. These core nodes are then connected to a multitude of edge networks.

An edge network connects one or more core nodes to, say, 10 to 100 edge nodes, so the number of edge nodes can be as many as a few thousand. An edge node can support thousands of war fighters, sensors, and unmanned systems. It uses links, carrying a moderate number of wavelengths, and forms the optical infrastructure supporting tactical wireless networks.

Today, the optical core of the DoD Enterprise Network is GIG-BE, the Global Information Grid Bandwidth Expansion, now managed by the Defense Information Systems Agency (DISA). It is a state-of-the-art, fiber-optic network with the impressive projected aggregate capacity of 10 terabits per second. Though it utilizes some advanced all-optical switching technology, it manages the traffic mostly by using optical-to-electrical-to-optical (O-E-O) switches and routers. Today, this is state-of-the-art networking technology.

To enable powerful network centric warfare applications, we project that the next-generation optical core network better have at least 10 times the capacity of GIG-BE (or an aggregate capacity of at least 100 terabits per second). Simply scaling up today's technology to meet these enormously great demands is not a viable answer. If we actually try to do this, the mostly O-E-O switching and routing nodes will become too expensive, too big, too power hungry, and way, way too hot.

Simply putting today's technology on steroids is not a solution. Our challenge is to find the technology needed to solve this problem. Moreover, as done in GIG-BE for economic reasons, this technology better be commercially viable, without sacrificing performance.

This is a tall order, but I believe the way to get there is through maximizing the use of all-optical technology and minimizing the need for O-E-O conversion, not only for switching, but for traffic grooming. Grooming is a networking technique used to produce the most efficient distribution of traffic from edge to core switch or router, and back to the edge again. It's like the hub system in commercial air travel today. What makes or breaks an airline is if it can keep all its planes, those from the edge to the hub and those going out again, as full as possible.

All-optical grooming (i.e., grooming without going through O-E-O conversion), is analogous to, but not as scary as, people changing planes in midair, without the planes having to land at a hub.

This model still requires electronic switches and routers to provide fine granularity circuit and packet grooming. But now, they will be cheaper, smaller, require less power, and generate less heat. Thus solving the three concerns I raised earlier associated with simply scaling-up today's technology.

In this vision, the optical layer deals only with switching and grooming of fast-configurable circuits. Packet switching and grooming is still done in the electronic domain. There is an ongoing DARPA program, DOD-N, or Data in the Optical Domain Networking, that investigates the feasibility of optical packet switching. The lessons learned and the technology generated in that program, will clearly be incorporated into our future thinking.

In fact, many DARPA programs over the years are helping to shape our current thinking and create the technologies for all-optical networking, dense wavelength-division multiplexing, and aggressive optoelectronic integration. Moreover, in spite of the recent downturn in the telecommunications industry, many other new ideas in optical networking have emerged. We need to make use of this rich set of new tools.

These are our 100-foot cedar trees. To me, chopping the cedars into 3-foot planks is like chopping a long optical path in the network into O-E-O sections only because we're more familiar with this technology. Even though this is my vision, I am happy to consider any other technology as long as it is reliable and meets all of the requirements I am outlining.

What are the requirements of the global DoD enterprise network of the future?

The first is huge capacity to enable real-time data fusion, visualization and synchronization, all providing universal situational awareness to the warfighter as well as to the armchair general.

The next-generation DoD enterprise network will be taking in sensor information from a variety of sources—satellites in space, manned and unmanned systems in the air, at sea and on the ground, soldiers in the field, and intelligence from a variety of places—all being transmitted to and from its edge nodes.

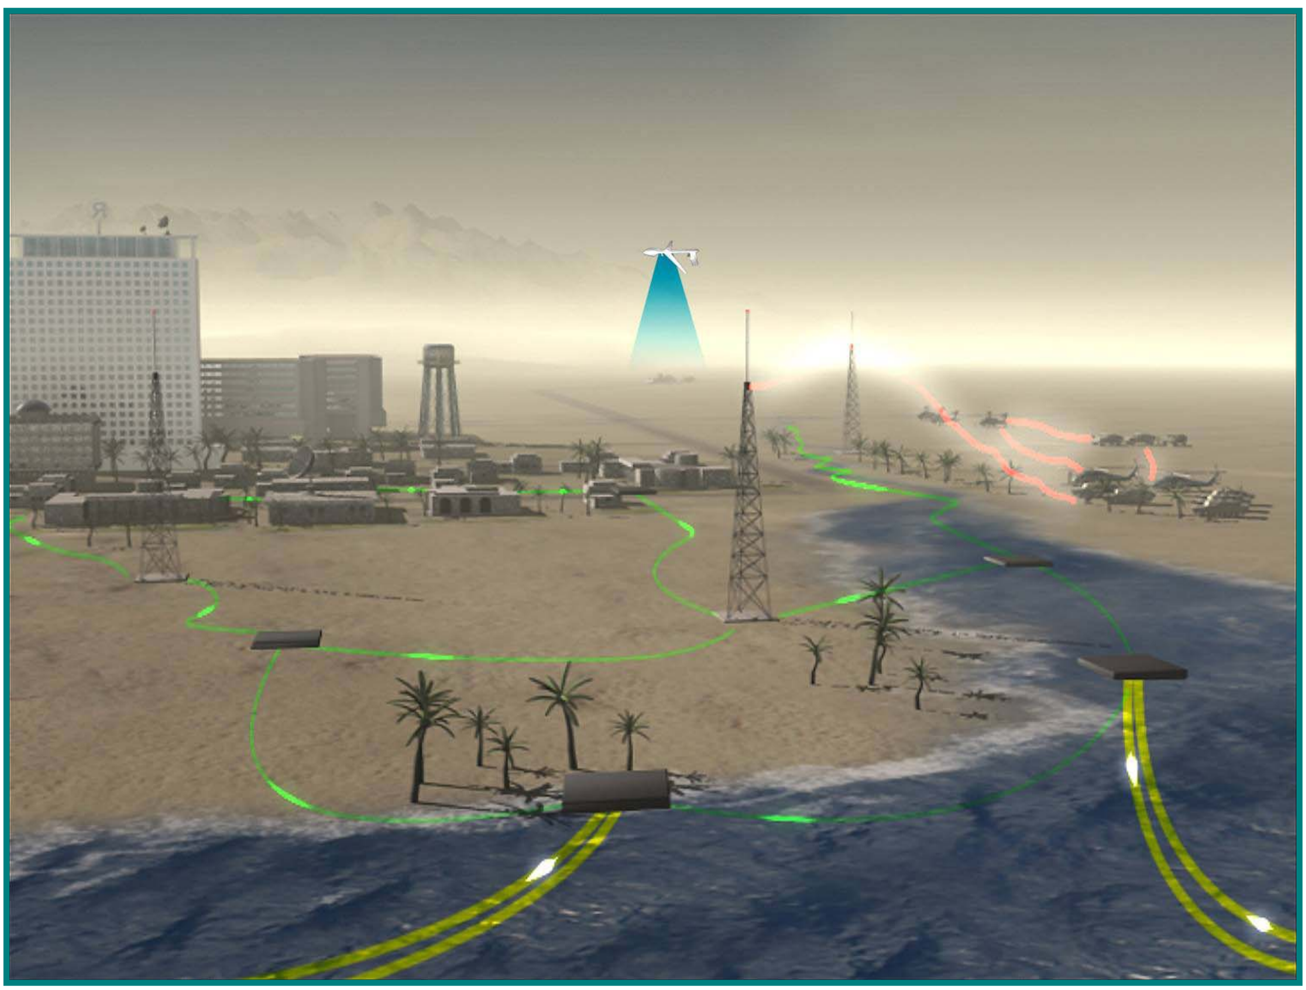

Next-Generation Global DoD Enterprise Network

Simultaneously, the network is invoking its memory, calling up huge databases and vast stores of knowledge. And, it is transmitting all of this to the various brains, the computers, which, in this case, may be distributed around the world.

Once there, the data is sorted and sifted, mined and fused, then transmitted back again to those who need it as intuitively actionable information, including 3D and even holographic images.

That is why we are saying that the capacity of the core network will have to be huge, on the order of 100 terabits per second; or 10 times today's state-of-the-art GIG-BE.

We have to be able to scale up to the required 100 terabits per second without simultaneously scaling up in cost, size, and power. We can't just build a bigger mousetrap, we have to build one that's many times as efficient. So that's the first requirement.

Second, it is often necessary for the network to distribute raw sensor data from multiple sources. This may require the transmission of different signal formats and protocols, which are naturally compatible with all-optical technology. Finally, the network needs fast configurability and low latency, to enable these interactive applications and provide very fast recovery from multiple failures.

Think of war conditions. The system has to be able to take several hits at once and re-route the traffic around the faults within a blink of an eye. With the real-time information demands of future networked warfare, we simply can't afford to have the information flow cut for the seconds or minutes—and sometimes much more—which is the time it takes for today's networks to recover from failures. The next-generation network has to work seamlessly in what one must assume will be a very hostile environment.

Again, I repeat, it's my belief that the way to get all that in a scalable way is through maximizing the use of all-optical technology.

The final challenge that I want to present to you, is that of an optical edge network with special

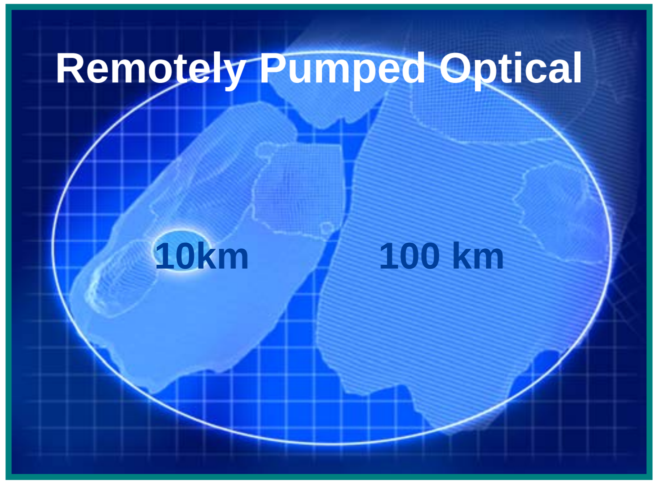

characteristics. In friendly areas, the edge network would resemble the core network I've just described, but with reduced capacity. Those parts deployed in unfriendly areas would be different. For example, it would be highly desirable for the switching nodes to be totally passive, i.e., they require no power. The optical switching nodes would be fixed, with the direction of the signal determined by the wavelength of the optical signal. In this case, one needs to use tunable transmitters and receivers at the edge of this passive network.

This passive networking technology can allow an edge network to grow to span tens of kilometers, but the edge network could be extended to hundreds of kilometers, using optical amplifiers that are remotely pumped from the edge to keep the signal from attenuating. In other words, you want to be able to just lay it in the ground or in riverbeds, hundreds of kilometers distant from the nearest power supply, and it will still work.

Most importantly, to meet the requirements for high reliability, this edge network has to be highly connected and self-healing so that it can sustain multiple failures and still perform without a hitch. One can also think of deploying an optical network within a plane or a ship to handle all of the communications needs in these platforms. Here, network reliability is still a top priority of course, but what can give the fiber-optic technology the edge over current copper-based technologies, are the increased capacity and reliability, the reduced weight, and the immunity to electromagnetic interference.

I realize that these are all tall orders, but you can't achieve the fantastic advances promised by truly network centric warfare without comparable advances in the thing that makes it all possible, the network itself.